

# Optimizing optical follow-up of gravitational-wave candidates

Leo Singer

leo.singer@ligo.org

Larry Price

larryp@caltech.edu

*LIGO Laboratory, California Institute of Technology  
MC 100-36, 1200 E. California Blvd., Pasadena, CA 91101, USA*

Antony Speranza

asperanz@mit.edu

*Department of Physics, Massachusetts Institute of Technology  
4-304, 77 Massachusetts Ave., Cambridge, MA 02139, USA*

## ABSTRACT

Observations with interferometric gravitational-wave detectors result in probability sky maps that are multimodal and spread over  $\gtrsim 10$ – $100$  deg<sup>2</sup>. We present a scheme for maximizing the probability of imaging optical counterparts to gravitational-wave transients given limited observing resources. Our framework is capable of coordinating many telescopes with different fields of view and limiting magnitudes. We present a case study comparing three different planning algorithms. We find that, with the network of telescopes that was used in the most recent joint LIGO–Virgo science run, a relatively straightforward coordinated approach doubles the detection efficiency relative to each telescope observing independently.

*Subject headings:* gravitational waves — telescopes — methods: numerical — gamma ray burst: general

## 1. Introduction

Among the most promising sources for co-incident electromagnetic (EM) and gravitational-wave (GW) emission are compact binary coalescence (CBC) events in which two neutron stars (NSs) or a NS and a black hole (BH) inspiral and merge. The NS may be torn apart before merger, powering EM emission at a variety of timescales and wavelengths ranging from seconds to months in  $\gamma$ - and X-ray to radio, respectively (Mészáros 2006; Nakar 2007; Metzger & Berger 2012). It is suspected that CBCs are also the progenitors (Belczynski et al. 2006; Lee & Ramirez-Ruiz 2007a; Troja et al. 2008) of some or all short, hard  $\gamma$ -ray bursts (SHGRBs). Plausible CBC event rates suggest that Advanced LIGO (aLIGO; Harry & the LIGO Scientific Collaboration 2010) and Advanced Virgo (AdVirgo; Virgo Collaboration

2009) could detect about 40 NS–NS and 10 NS–BH events per year of observation time (Abadie et al. 2010). Aside from shedding light on the true progenitors of SHGRBs (Sylvestre 2003; Lee & Ramirez-Ruiz 2007b), coincident detections of EM and GW emission could also be utilized as “standard sirens” to provide a precision measurement of the Hubble constant (Schutz 2001; Nissanke et al. 2010).

In addition to CBCs, there are other possible sources of coincident EM and GW emission. Examples include type II supernovae, which are much more common than CBCs at  $\lesssim 1$  event per 20 years in the Milky Way alone, but whose GW counterparts would be detectable by Advanced LIGO (aLIGO) only within the local group (Ott 2009). Soft  $\gamma$  repeaters (Ioka 2001; Abbott et al. 2008) and anomalous X-ray pulsars (Abadie et al.

2011) could produce GWs, and some models predict GW emission comparable to the best upper limits (Corsi & Owen 2011). Neutron stars recovering from pulsar glitches might produce GW transients on timescales of days to weeks that could be detectable at  $\lesssim 10$  kpc (Bennett et al. 2010). More speculative sources of GW include cosmic string cusps (Damour & Vilenkin 2000; Siemens et al. 2006), which if they exist might also power GRBs (Babul et al. 1987; Berezhinsky et al. 2001; Wang et al. 2011) and radio bursts (Vachaspati 2008).

The advent of deep, wide-field optical survey telescopes and rapid response robotic telescopes is paving the way for time-domain astronomy. The simultaneous construction of advanced GW detectors creates the enticing possibility of a coincident detection of a GW transient and an associated EM counterpart. Efforts in this direction were recently made in LIGO’s sixth and Virgo’s second and third science runs, which saw the introduction of a program (Kanner et al. 2008; Abadie et al. 2012b,a; Aasi et al. 2012b,a) whereby potential GW transient alerts were sent to X-ray, optical, and radio facilities for follow-up. In the context of space-based (Kocsis et al. 2008) and advanced ground-based (Fairhurst 2009; Cannon et al. 2012; Manzotti & Dietz 2012) GW detectors, there is even discussion of using GWs to trigger early-warning follow-up observations while the inspiral is still ongoing.

The great disparity in the character of optical and GW observations makes GW-triggered detection of optical counterparts challenging. Although there are now several operating ultra-wide-field, but relatively shallow, transient monitoring telescopes such as Pi of the Sky, only a handful of operating or planned deep survey telescopes have fields of view (FOVs) that are comparable to the error regions of  $\sim 10$ – $100$  deg<sup>2</sup> (Fairhurst 2009; Wen & Chen 2010; Nissanke et al. 2011; Fairhurst 2011) that will be typical of GW sky maps in the advanced detector era. By contrast, a GW detector is both blessed and cursed by its nature as an all-sky instrument: it has coarse directional sensitivity only due to its geometry-defined quadrupole antenna pattern.

In this paper we lay a framework for planning optimal EM follow-up of GW candidates. Our paradigm focuses on tiling the GW sky map

as efficiently as possible by coordinating a number of telescopes with dissimilar FOVs. Although our formalism supports it, we do not yet account for seeing, different limiting magnitudes, different times at which observations are made, or multiple observations with each telescope. As preliminaries, we describe the key features of ground-based interferometric GW detectors as contrasted with optical telescopes. We demonstrate that the optimal pointing of a single telescope can be expressed as a convolution-like integral on the unit sphere, which is amenable to spherical harmonic analysis. For multiple telescopes, we provide a selection of observation planning algorithms that take into account the particular geometry of each telescope’s image plane as well as local observing conditions (in our present implementation, just the position of the sun and the Earth). We have developed a parallelized observation planning code in Python/C/C++ called BAYEsian optimal Search for Transients with Autonomous and Robotic telescopes (BAYESTAR)<sup>1,2</sup>, which can map out follow-up telescope pointings with low latency in the way described in this work. We show that with the telescopes that were available for EM follow-up in the last LIGO–Virgo science run, if each telescope performs just one pointing, the imaging efficiency is doubled by pointing all of the telescopes in coordination rather than independently. Finally, we propose an expanded simulation campaign and discuss some practical issues for EM follow-up in the advanced GW detector era.

## 2. GW detectors and optical telescopes

In this section, we review the observational capabilities of GW detectors as contrasted with optical telescopes.

GW detectors such as LIGO (Abbott et al. 2009) and Virgo (Acernese et al. 2008) are Fabry–Perot Michelson laser interferometers that

<sup>1</sup>A pun on the Cylon battleships in the American television series *Battlestar Galactica*. The defining characteristic of the Cylons is that they repeatedly defeat humanity by using their superhuman information-gathering ability to coordinate overwhelming forces. The name also suggests that, like the Cylons, robotic transient telescopes may some day rise against us humans.

<sup>2</sup>We do not like to mention the final ‘T’ in the acronym, because then it would be called BAYES-TART, which would sound ridiculous.

sense GW-induced strain through differences in light travel time in two arms. The only directional sensitivity of a single GW detector is due to its antenna pattern (see, for example, [Sathyaprakash & Schutz 2009](#)). As predicted in general relativity, GWs are transverse and come in two orthogonal polarizations, the so-called ‘+’ and ‘×’ modes. An interferometer with perpendicular arms is most sensitive to a ‘+’ polarized source at the zenith (or nadir), with the sensitivity going to zero at any azimuth that is  $45^\circ$  from either arm. The sensitivity to the ‘×’ mode is also greatest at the zenith, but vanishes at any azimuth that is at a right angle to either arm. It also vanishes everywhere in the plane of the detector.

A signal’s time delay on arrival at two GW detectors constrains the position of the GW source to an annulus on the sky. The antenna patterns of the two detectors impose constraints that disfavor parts of the annulus ([Raymond et al. 2009](#)). Time delay measurements from three gravitational wave detectors constrain a source to two regions that are located at mirroring points on either side of the plane formed by the detectors. The mirror degeneracy is, again, partly broken by the antenna patterns. Networks of four or more detectors are sufficient to determine the location of the source to within a single patch on the sky. The network consisting of Advanced Virgo (AdVirgo) and the two aLIGO observatories will be able to constrain sources to areas of  $10\text{--}100\text{ deg}^2$  with 90% confidence, varying across the sky as a result of the antenna patterns ([Wen & Chen 2010](#); [Nissanke et al. 2011](#); [Fairhurst 2011](#); [Klimenko et al. 2011](#)). Adding the planned Japanese detector KAGRA (formerly LCGT; [Kuroda 2010](#)) or the proposed LIGO–India detector would narrow the confidence region ([Fairhurst 2011](#); [Schutz 2011](#)).

Optical telescopes are markedly different from gravitational wave detectors in that they have directional precision of arcseconds or better. This is at the expense of observing only a small fraction of the sky at one time. Table 1 shows a list of telescopes that participated in the LIGO–Virgo EM follow-up program. FOVs vary from scarcely  $0.01\text{ deg}^2$  with relatively deep instruments such as Liverpool Telescope at a limiting 21 mag; to a few  $\text{deg}^2$  with small robotic transient follow-up telescopes such as ROTSE at 17 mag; to many  $\text{deg}^2$

with survey instruments such as Skymapper and Palomar Transient Factory (PTF), at  $\gtrsim 20$  mag. The largest FOVs of  $\gtrsim 100\text{ deg}^2$  belong to transient monitoring instruments such as Pi of the Sky at  $\gtrsim 12$  mag. The eagerly awaited Large Synoptic Survey Telescope (LSST) will have an even larger FOV than PTF, but with a limiting magnitude of 24.5 in a similar exposure time. Slew times can be as brief as seconds for small robotic telescopes like ROTSE and TAROT, or minutes and longer for larger instruments or instruments that are less optimized for response time. Typical total exposure times range from  $\sim 10$  s to a few minutes per pointing.

### 3. Single telescope case

In this section, we introduce the optimization problem of pointing a single telescope to maximize the probability of imaging the true, but unknown, location of the GW source. Then, we describe two numerical procedures to find the optimum.

For GW detectors, the observables are strain time series consisting of noise plus the astrophysical GW signal. The GW signal from a CBC source is determined by extrinsic parameters consisting of right ascension  $\phi$ , declination  $(\pi/2 - \theta)$ , luminosity distance  $D_L$ , inclination angle  $\iota$ , coalescence time  $t$ , polarization angle  $\psi$ , and orbital phase at coalescence  $\varphi$ ; and intrinsic parameters including the component masses and spins. One may form the posterior distribution of all of these parameters:

$$p(\theta, \phi, D_L, \iota, t, \psi, \varphi, m_1, m_2, \dots | \text{GW}).$$

The proposition GW represents the entire GW observation: the strain in each GW detector, the geographic locations of the detectors, and their antenna patterns. If one fixes or marginalizes over all parameters except for the position on the sky, the spherical polar coordinates  $\omega = (\theta, \phi)$ , then this reduces to what we term the GW sky map,

$$p(\omega | \text{GW}) = s(\omega).$$

This quantity can be evaluated from time delays and antenna factors ([Abadie et al. 2012a](#)) or from the GW strain itself using analytic methods ([Searle et al. 2008, 2009](#)) or Markov chain Monte Carlo (MCMC) techniques ([Christensen & Meyer 1998](#); [van der Sluis et al. 2008](#); [Veitch & Vecchio 2010](#); [Nissanke et al. 2011](#)).

TABLE 1  
TELESCOPES EMPLOYED IN JOINT LIGO–VIRGO SCIENCE RUN

site	field of view	limiting magnitude	slew time (s)	geographic location
Liverpool <sup>ab1</sup>	$0.077^\circ \times 0.077^\circ$	21 in 300 s	30	$28^\circ 45' 44.8''$ N, $17^\circ 52' 45.2''$ W
Zadko <sup>c</sup>	$1.4^\circ \times 1.4^\circ$	21 in 180 s	20	$31^\circ 21' 24''$ S, $155^\circ 42' 49''$ E
ROTSE III-a <sup>d</sup>	$1.85^\circ \times 1.85^\circ$	17 in 5 s	4	$31^\circ 16' 24.1''$ S, $149^\circ 3' 40.3''$ E
ROTSE III-b <sup>d</sup>	$1.85^\circ \times 1.85^\circ$	17 in 5 s	4	$23^\circ 16' 18''$ S, $16^\circ 30' 00''$ E
ROTSE III-c <sup>d</sup>	$1.85^\circ \times 1.85^\circ$	17 in 5 s	4	$36^\circ 49' 30''$ N, $30^\circ 20' 0''$ E
ROTSE III-d <sup>d</sup>	$1.85^\circ \times 1.85^\circ$	17 in 5 s	4	$30^\circ 40' 17.7''$ N, $104^\circ 1' 20.1''$ W
TAROT <sup>ef</sup>	$1.86^\circ \times 1.86^\circ$	17 in 10 s	1.5	$43^\circ 45' 8''$ N, $6^\circ 55' 26''$ E
TAROT-S <sup>ef</sup>	$1.86^\circ \times 1.86^\circ$	17 in 10 s	1.5	$29^\circ 15' 39''$ S, $70^\circ 43' 56''$ W
Skymapper <sup>g</sup>	$2.373^\circ \times 2.395^\circ$	21.6 in 110 s		$31^\circ 16' 24''$ S, $149^\circ 3' 52''$ E
PTF <sup>h</sup>	$3.5^\circ \times 2.31^\circ$	20.6 in 60 s		$33^\circ 21' 21''$ N, $116^\circ 51' 50''$ W
QUEST <sup>i</sup>	$3.6^\circ \times 4.6^\circ$	21 in 140 s		$33^\circ 21' 21''$ N, $116^\circ 51' 50''$ W
Pi of the Sky South <sup>jk</sup>	$20^\circ \times 20^\circ$	12 in 10 s	60	$22^\circ 57' 12''$ S, $68^\circ 10' 48''$ W
Pi of the Sky North <sup>jk</sup>	$40^\circ \times 40^\circ$	12 in 10 s	40	$37^\circ 6' 14''$ N, $6^\circ 44' 3''$ W

NOTE.—FOV, limiting magnitude, slew time, and geographic location of ground-based robotic optical telescopes that participated in the LIGO–Virgo EM follow-up program. Except where otherwise indicated, we quote the shortest exposure for which a limiting magnitude is available in the literature.

<sup>1</sup>LIGO’s request for Liverpool Telescope time calls for a 300 s exposure based on a limiting magnitude of 21 (White 2012).

References. — <sup>a</sup>Gomboc et al. (2004); <sup>b</sup>Guidorzi et al. (2006); <sup>c</sup>Coward et al. (2010); <sup>d</sup>Akerlof et al. (2003); <sup>e</sup>Boër et al. (2003); <sup>f</sup>Klotz et al. (2008); <sup>g</sup>Keller et al. (2007); <sup>h</sup>Law et al. (2009); <sup>i</sup>Baltay et al. (2007); <sup>j</sup>Małek et al. (2010); <sup>k</sup>Majcher et al. (2011);

Consider the  $i$ th of  $N$  telescopes labeled  $1 \dots N$ . Let the proposition  $\text{EM}_i$  denote the detection of an associated optical transient with telescope  $i$ . Let  $\mathbf{C}_i$  represent the configuration of telescope  $i$ , which consists of a pointing  $\gamma_i = (\theta_i, \phi_i)$ , nominal limiting magnitude, exposure start and end times, as well as uncontrollable factors such as seeing and interference from the horizon, sun, moon, and the galactic disk. If we knew the position  $\omega$  of the source, we could evaluate the probability of detecting an EM counterpart at  $\omega$  with telescope  $i$ ,

$$p(\text{EM}_i | \mathbf{C}_i, \omega).$$

For most of this paper we shall restrict ourselves to the dependence on the pointing of the telescope, so that the detection probability factors into a function  $b_i(\omega)$  describing the telescope's FOV in a home orientation (i.e., for equatorial mounts, the celestial North pole) and a function  $v_i(\omega)$  that describes environmental factors including the night-time/twilight terminator and the horizon. In principal,  $v_i(\omega)$  could also include other observational factors such as moonlight and current seeing. The detection probability becomes:

$$p(\text{EM}_i | \gamma_i, \omega) = b_i(R^{-1}[\gamma_i] \omega) v_i(\omega). \quad (1)$$

A simple rectangular FOV with angular width  $w$  and height  $h$  is described by

$$b_i(\theta, \phi) = \begin{cases} 1 & \text{if } |\sin \theta \cos \phi| \leq \sin w/2, \\ & |\sin \theta \sin \phi| \leq \sin h/2, \\ & \text{and } \theta < \pi/2; \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Much more complicated FOVs are commonplace because a real telescope may have a focal plane consisting of a CCD mosaic with a filling factor that is much less than unity; dead CCDs or CCDs that contain defects; or CCDs that are vignetted or clipped by the optics.

$R[\gamma_i]$  denotes an active rotation operation about the  $z$ -axis by  $\phi_i$  and about the  $y$ -axis by  $\theta_i$ , so  $R^{-1}[\gamma_i]$  is the passive rotation such that  $b_i$  evaluated at  $(R^{-1}[\gamma_i] \omega)$  represents the FOV centered on the sky location  $\gamma_i = (\theta_i, \phi_i)$ . This rotation matrix is described by two, not three, angles because mounts for ground-based telescopes usually have two degrees of freedom (for example, right ascension and declination or azimuth and altitude).

This particular model incorporates the assumption that a telescope always detects an optical transient when the true location of the GW source is inside the FOV. We can identify this with the probability of *imaging* the source, which for most of this work we shall use as a proxy for the ultimately desirable probability of *detection*.

The optimal configuration  $\mathbf{C}_i^*$  of telescope  $i$  maximizes the probability of detection of an electromagnetic counterpart given all gravitational wave observations, marginalized over the unknown source location. To wit,

$$\mathbf{C}_i^* = \arg \max_{\mathbf{C}_i} \int p(\text{EM}_i | \mathbf{C}_i, \omega) p(\omega | \text{GW}) d\Omega, \quad (3)$$

where the integral is over  $\omega$  and  $d\Omega$  represents the surface element on the unit sphere. In the simplified scenario in which the pointing orientation is the only free parameter, the optimal pointing is given by

$$\begin{aligned} \gamma_i^* &= \arg \max_{\gamma_i} \int p(\text{EM}_i | \gamma_i, \omega) p(\omega | \text{GW}) d\Omega \\ &= \arg \max_{\gamma_i} \int s(\omega) b_i(R^{-1}[\gamma_i] \omega) v_i(\omega) d\Omega. \end{aligned} \quad (4)$$

The integral in Equation (4) resembles a convolution integral, but over rotations instead of over the real number line. Just as fast Fourier transform (FFT) methods are used to efficiently evaluate convolutions, fast spherical harmonic methods may be used to evaluate Equation (4) efficiently. In the Appendix, we describe in detail two algorithms for solving the single telescope problem: a “direct” or spatial algorithm and a “multipole” algorithm based on the method for fast convolution on the sphere of Wandelt & Górski (2001). Both algorithms exploit special properties of the HEALPix<sup>3</sup> (Górski et al. 2005) projection. Either of these algorithms may be used as a component of the planning of observations with multiple telescopes, which we discuss in the next section.

#### 4. Multiple telescope case

In order to address the problem of pointing multiple telescopes, we must first specify a figure of

<sup>3</sup><http://healpix.jpl.nasa.gov/>

merit. One example is the probability of observing an electromagnetic counterpart in any (one or more) telescope,

$$p_{\text{any}} \equiv p(\text{EM}_1 \cup \text{EM}_2 \cup \dots \cup \text{EM}_N | \gamma_1, \gamma_2, \dots, \gamma_N, \text{GW}). \quad (5)$$

This only considers the case in which each telescope is asked to point once. Some telescopes will take a sequence of images at arbitrary sky locations. Others are constrained to follow a particular path on the sky, for instance, visiting reference pointings in a specified order. In the former case, we could trivially interpret  $\text{EM}_1 \dots \text{EM}_N$  as optical detections in the *observations*  $1 \dots N$  at times  $t_1 \dots t_N$  respectively, where there are  $N$  observations but potentially  $\ll N$  telescopes. In the latter case, we could define a composite FOV that describes the probability of imaging the source in any of the pointings throughout the entire observation. A more elaborate objective function would also include a penalty on the number of pointings in order to economize the limited observational resources. For the remainder of this paper, we will focus on  $p_{\text{any}}$ <sup>4</sup>.

De Morgan's law tells us that Equation (5) is equivalent to the complement of the probability of

<sup>4</sup> Alternatively, we could demand not one, but several, EM detections to secure the discovery of a bona fide optical transient. Combinatorial considerations give the probability of detection in at least  $K$  out of  $N$  observations,

$$\begin{aligned} p(\text{at least } K \text{ detections} | \text{GW}) = & \int \left[ (-1)^{K+1} \binom{0}{K-1} \sum_i p(\text{EM}_i | \gamma_i, \omega) \right. \\ & + (-1)^{K+2} \binom{1}{K-1} \sum_{i \neq j} p(\text{EM}_i \cap \text{EM}_j | \gamma_i, \gamma_j, \omega) \\ & + (-1)^{K+3} \binom{2}{K-1} \sum_{i \neq j \neq k} p(\text{EM}_i \cap \text{EM}_j \cap \text{EM}_k | \gamma_i, \gamma_j, \gamma_k, \omega) \\ & + \dots \\ & \left. + (-1)^{K+N} \binom{N-1}{K-1} p(\text{EM}_1 \cap \text{EM}_N | \gamma_1, \dots, \gamma_N, \omega) \right] \\ & p(\omega | \text{GW}) d\Omega. \end{aligned} \quad (6)$$

non-detection with every telescope,

$$\begin{aligned} p_{\text{any}} &= 1 - \int \left( \prod_i [1 - p(\text{EM}_i | \gamma_i, \omega)] \right) p(\omega | \text{GW}) d\Omega \\ &= 1 - \int [1 - b_1(R^{-1}[\gamma_1] \omega) v_1(\omega)] \\ &\quad [1 - b_2(R^{-1}[\gamma_2] \omega) v_2(\omega)] \\ &\quad \dots [1 - b_N(R^{-1}[\gamma_N] \omega) v_N(\omega)] s(\omega) d\Omega. \end{aligned} \quad (7)$$

The optimal set of pointings maximizes  $p_{\text{any}}$ ,

$$(\gamma_1^*, \gamma_2^*, \dots, \gamma_N^*) \equiv \arg \max_{\gamma_1, \gamma_2, \dots, \gamma_N} p_{\text{any}}. \quad (8)$$

#### 4.1. Planning algorithms

We considered three different multiple telescope algorithms, increasing in sophistication and accuracy but also increasing in computational overhead. After describing each algorithm in detail, we exhibit implementations that use the positional astronomy library NOVAS<sup>5</sup> (Barron et al. 2011) to take into account rudimentary sun and horizon interference assuming observers at sea level on an atmosphereless, perfectly spherical Earth. We assume that all observations occur simultaneously at the time of the GW detection.

##### 4.1.1. Noncoordinated

The *noncoordinated* planner models the situation in which each telescope operator independently decides to observe that patch of the sky that maximizes his or her own odds of detecting a counterpart. Each single telescope problem is solved independently, regardless of the configurations of any of the other telescopes.

The *noncoordinated* method has the disadvantage that it will generally result in all of the telescopes pointing at the same region of the sky, perhaps failing to image all of the significant regions of the GW sky map. However, if the telescope network's coverage is exceedingly poor, such that the nonzero regions of the sky map are within view of only  $\sim 1$  telescope at a time, then this will not be a significant drawback. The *noncoordinated* method

<sup>5</sup> [http://aa.usno.navy.mil/software/novas/novas\\_info.php](http://aa.usno.navy.mil/software/novas/novas_info.php)



also has the advantage that it is exceedingly simple to implement. For these reasons, the *noncoordinated* planner may represent a reasonable, but sub-optimal, observational strategy.

#### 4.1.2. Greedy

The *greedy* planner is slightly more sophisticated than the *noncoordinated* method. Suppose that pointings have been decided upon for telescopes  $1, 2, \dots, i$ . Then the configuration of telescope  $i + 1$  is chosen such that it maximizes the imaging probability when all of the previous pointings  $1 \dots i$  are held fixed. This procedure is illustrated in Figure 1. It is not guaranteed to find an absolute maximum, but we shall see that in practice it does quite well.

This algorithm leaves a choice of the order in which to point the telescopes. This is analogous to the problem of packing a number of stones of different sizes into a bucket. The intuitive answer is to insert the largest rocks first and then use the pebbles to fill in the gaps. Likewise, we find that we obtain the best detection efficiency when we sort the telescopes by decreasing FOV area because the telescope with the largest FOV makes the greatest progress toward imaging the entire support of the GW sky map. We refer to the greedy algorithm with this prescription as the *greedy-sorted* planner. We expect that the *greedy-sorted* planner will always outperform the *noncoordinated* planner by a factor that depends on the signal-to-noise ratio (SNR) of the GW signal and the particular network of telescopes.

#### 4.1.3. Simulated annealing

The *anneal* planner uses the simulated annealing code provided by the SciPy library<sup>6</sup> to extremize  $p_{\text{any}}$ . Starting from an initially proposed observing plan, this procedure randomly perturbs all of the telescopes' pointings simultaneously. If the imaging probability has improved, the perturbed plan is adopted. The perturbations become smaller and smaller with each iteration; the perturbation distribution as a function of time is called a *cooling schedule*. We used a modified version of the “very fast” cooling schedule described by Ingber (1989) with an initial “temper-




Fig. 1. Flow chart depicting greedy algorithm.

<sup>6</sup><http://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.anneal.html>

ature”  $T_0 = 10^{-4}$ . Since [Ingber](#)’s cooling schedule uses rectilinear coordinates, we modified it by perturbing the Cartesian coordinates of the center of each telescope’s FOV and then normalizing to unit length.

Simulated annealing will nearly always find the globally optimal observing plan in a finite, but possibly large, number of iterations. With a particular fixed number of iterations, we expect that it will usually outperform not only the *noncoordinated* planner but also the *greedy-sorted* planner. However, the computational overhead of evaluating thousands of observing plans in sequence may be prohibitive if there is a time constraint—if, for example, it is important to begin taking images as soon as possible after the GW event.

## 4.2. Case study

We compared the imaging efficiency of our three multiple telescope planners using a set of 2126 GW sky maps from signals injected into simulated initial LIGO and Virgo noise over a period of 24 hours. The simulated GW signals are accurate to second post-Newtonian order in phase. Component masses are drawn from a uniform distribution from 1 to 15  $M_\odot$  with total mass is restricted to  $<20 M_\odot$ . Distances are drawn randomly from a log-uniform distribution from 10 to 40 Mpc. A sky map was produced for each injection using the time delay method described in [Abadie et al. \(2012a\)](#). We generated two observing plans for each sky map for each of the three planners. The first observing plan used  $v_i(\omega) = 1$  for all observations  $i$  and sky locations  $\omega$ , disregarding the positions of the sun and Earth. The second observing plan used  $v_i(\omega) = 1$  above the horizon and at all sky locations from which the sun was at least  $18^\circ$  below the horizon, and zero elsewhere, simulating observing conditions on an atmosphereless, spherical Earth. The output of the *greedy-sorted* planner was used as the initial proposal for simulated annealing. Pi of the Sky and Liverpool Telescope were excluded from the simulation because the areas of their FOVs are outliers with regard to the other FOVs of  $\sim 1\text{--}10 \text{ deg}^2$  represented in the model. Figure 4 illustrates an observing plan generated with the greedy algorithm with sun and horizon constraints disabled.

Figure 2 shows the measured imaging efficiency as a function of the luminosity distance. For each

sky map and for each planner, we tested whether the injected location would have been imaged, i.e., whether the injected source location was inside any of the FOVs of the telescopes. Dividing the injection set into six concentric shells, each 5 Mpc deep, from 10 to 40 Mpc, we tallied the number of sources that would have been imaged and the total number of sources. Because the injections were drawn from a uniform distribution in  $\log D_L$ , we weighted each injection by

$$D_L^2 \left| \frac{\partial \log D_L}{\partial D_L} \right|^{-1} \equiv D_L^3 \quad (9)$$

in order to approximate a uniform-in-volume distribution. The dashed lines represent the observing plans that disregard the presence of the sun and the Earth whereas the solid lines represent the observing plans that fully account for these effects.

For each injection, we also computed the probability  $p_{\text{any}}$  of imaging the source, marginalized over  $(\theta, \phi)$  and not making use of knowledge of the true sky location, using Equation (7). We then formed a histogram of  $p_{\text{any}}$  over all of our injections, again weighting by the correction factor  $D_L^3$ . This histogram is shown in Figure 3.

Two systematic effects are apparent in Figure 2. First, for all planning algorithms, accounting for the sun results in a reduction in efficiency by a factor of  $\approx 1/3$ , independent of distance. This is because astronomical twilight begins and ends when the sun is  $18^\circ$  below the horizon ([Astronomical Almanac 2010](#)); solar illumination is faint enough for astronomical observations from the Earth’s surface over over a fraction  $(1 - \cos[90^\circ - 18^\circ]) / 2 \approx 0.35$  of the sky. Second, the efficiency falls as luminosity distance increases because the area of the GW error region grows as signal-to-noise ratio decreases.

Both Figures 2 and 3 demonstrate that co-ordination of many optical telescopes increases the chances of imaging the true location of the GW source. Whether or not sun and horizon interference are included, the *anneal* and *greedy-sorted* planners achieve about twice the detection efficiency of the *noncoordinated* planner for the network of telescopes used in the last joint LIGO–Virgo science run and in our model.

From Figure 3, we see that simulated annealing generally outperforms the greedy algorithm.






Fig. 2.— Average detection efficiency of *noncoordinated*, *greedy-sorted*, and *anneal* planners, for sources located within spherical shells 5 Mpc thick at 10–40 Mpc. Solid lines include sun and horizon interference; dashed lines ignore interference considerations.




Fig. 3.— Histogram of expected detection efficiency  $p_{\text{any}}$  as computed using Equation (7), for sources uniformly distributed in volume from 10 to 40 Mpc. As in Figure 2, solid lines include sun and horizon interference and dotted lines ignore interference considerations.

At least a slight improvement is expected because the observing plan generated by the *greedy-sorted* algorithm is used as the initial proposal for simulated annealing. Simulated annealing would also give only a slight improvement if the *greedy-sorted* observing plan was already nearly optimal.

## 5. Discussion

We have presented a framework for planning optimal EM follow-up of GW transients. We have shown that, if each telescope images one sky location, a coordinated approach with a particular network of telescopes is twice as likely to image the true location of a GW event, versus an uncoordinated approach. It remains to be seen whether coordination will be important for more realistic observing strategies, which will involve taking multiple images with each telescope. A successful observational program also will likely involve several different stages calling for markedly different kinds of instruments. In the context of a grand strategy for the follow-up of GW candidates (Metzger & Berger 2012), spatially coordinated observing campaigns like those we have considered are most likely to be applicable to prompt, rapidly fading, on-axis afterglows, though it is unclear whether coordination will be relevant for later-rising and fainter counterparts. This will be the subject of future work.

We have found that a simple greedy algorithm performs almost as well as simulated annealing while taking much less time on average. In the interest of low latency, we have provided a parallel implementation of the convolution algorithm so that observation plans can be generated rapidly on a multicore workstation. The observation planning code has been operated in a low latency mode during software engineering runs of the joint LIGO–Virgo low latency data analysis.

In this work, the objective function that we have used is the probability of *imaging* the source. This is a proxy for the more desirable objective function of the probability of *detecting* the source, which must take into account the limiting magnitudes of the observations, a description of the potential false positives, and a model of the EM emission. A weighting could be assigned to each telescope’s FOV based upon the integrated flux that the telescope would be able to gather given

the anticipated power law index of the light curve and the telescope's limiting magnitude, slew time, and exposure time. For example, assuming a power-law light curve  $F_\nu \propto t^{-\alpha}$ , at time  $t$  a telescope with limiting magnitude  $M$  would be able to detect a source up to a maximum luminosity distance  $D_{L,\max} \propto \sqrt{2.512^M t^{-\alpha}}$ . We can marginalize over the unknown distance to the source, assuming that nearby  $\gamma$ -ray bursts (GRBs) are uniformly distributed in space. Neglecting cosmological corrections, the detection probability of a single telescope might be expressed by

$$p(\text{EM}_i | \gamma_i, \omega) \propto \begin{cases} (2.512^{M_i} t_i^{-\alpha})^{3/2} & \text{inside FOV,} \\ 0 & \text{outside FOV,} \end{cases} \quad (10)$$

where  $t_i$  is the time of observation  $i$  and  $M_i$  is the limiting magnitude of observation  $i$ . We could also marginalize over the power-law index  $\alpha$  and the initial luminosity.

Although this prescription will be appropriate for a single telescope, for multiple telescopes or multiple pointings of the same telescope it will be more appropriate to form the probability of detection with the network of telescopes, conditioned on the luminosity distance, and *then* marginalize over luminosity distance. Cast in this form, the multiple telescope detection probability is

$$p_{\text{any}} = 1 - \frac{1}{V} \int \prod_i \left[ 1 - \Theta[D_{L,\max,i}(L_i, M_i) - D_L] b_i(R^{-1}[\gamma_i] \omega) v_i(\omega) \right] s(\omega) \\ p(L_1, L_2, \dots, L_N) dL_1 dL_2 \dots dL_N D_L^2 dD_L d\Omega \quad (11)$$

where  $V$  is the volume of integration and  $\Theta$  is the Heaviside step function.  $D_{L,\max,i} \propto \sqrt{2.512^{M_i} L_i}$  is the farthest detectable luminosity distance for observation  $i$ , being a function of the source's luminosity  $L_i$  at the time  $t_i$  of the observation and the limiting magnitude  $M_i$ . This figure of merit is not as simple as Equation (7), but it may behave better under simulated annealing because the FOVs no longer have to move rigidly around each other without overlapping. This is also the subject of future work.

Source code for BAYESTAR is available on the LIGO Data Analysis Software Working Group web site at <http://www.lsc-group.phys.uwm.edu/daswg/projects/bayestar.html>.

Some of the results in this paper have been derived using HEALPix (Górski et al. 2005).

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation (NSF) and operates under cooperative agreement PHY-0107417. Some results were produced on the NEMO computing cluster operated by the Center for Gravitation and Cosmology at University of Wisconsin–Milwaukee under NSF Grants PHY-0923409 and PHY-0600953. This research is supported by the NSF through a Graduate Research Fellowship to L.S. A.S. acknowledges support from the LIGO Laboratory Summer Undergraduate Research Fellowship, Visitors, and NSF Research Experiences for Undergraduates (REU) programs. This paper has LIGO Document Number LIGO-P1200033-v4.




Fig. 4.— Demonstration of a BAYESTAR observation plan for a simulated NS–NS CBC event injected into synthetic initial LIGO noise. The event produced triggers in the H1, L1, and V1 detectors. The GW sky map,  $p(\omega|\text{GW})$ , in posterior probability per square degree, is shown in color in the online version and in shades of gray in the print version. The FOVs of the telescopes are shown as shaded gray rectangles. Panels (a) and (b) show Lambert azimuthal equal area projections of two  $30^\circ \times 30^\circ$  patches where observations were planned. Panel (c) provides the selected telescope pointings in tabular form. Panel (d) shows the selected observations on an all-sky Mollweide projection.

## A. Implementation

Equation (4) is reminiscent of a convolution integral. This leads us to think of  $b_i(\boldsymbol{\omega})$  as an analogue of a finite impulse response (FIR) filter, but on the unit sphere. A FIR filter with a sufficiently compact kernel can be efficiently evaluated in so-called direct form by advancing the kernel one sample point at a time and taking an inner product with the signal. On the other hand, a FIR filter with a long, extended kernel is generally more efficiently implemented using frequency domain techniques such as the classic overlap-save algorithm (Numerical Recipes, Press et al. 2007), which exploits the convolution theorem and the FFT.

Similarly, if the telescope’s FOV spans sufficiently little area, then a discretized form of Equation (4) can be efficiently evaluated by subdividing both the FOV and the GW sky map into pixels, and for each pointing  $\gamma_i$ , rotating the pixels in the kernel and then taking their inner product with the pixels in the GW sky map. Transforming just the nonzero pixels of the kernel results in a substantial computational speedup; for a  $\sim 1 \text{ deg}^2$  FOV, this reduces the number of calculations by a factor  $4(180)^2/\pi \sim 4 \times 10^4$ .

If, however, the telescope’s FOV subtends a large solid angle, then the “direct” algorithm outlined above will be computationally expensive, with a cost growing as  $\mathcal{O}(n_{\text{pix}}^2)$ , where  $n_{\text{pix}}$  is the number of pixels covering the entire unit sphere. The FFT-based frequency domain techniques for FIR filters also have analogous fast *multipole* transform algorithms for functions defined on  $S^2$ . Considerable study has been devoted to fast multipole methods because of their utility in solving differential equations with approximate spherical symmetry. Example applications can be found in fields as varied as cosmic microwave background (CMB) mapping (Górski et al. 2005), numerical astrophysics, quantum chemistry, and even the entertainment industry (McEwen & Wiaux 2011). Presently, the best known upper bound on the complexity of both spherical harmonic transforms and fast convolution on the sphere (Górski et al. 2005; Wandelt & Górski 2001) is  $\mathcal{O}(n_{\text{pix}}^{3/2}) \sim \mathcal{O}(L^3)$ , where  $L \sim \mathcal{O}(n_{\text{pix}}^{1/2})$  is the azimuthal index of the highest non-vanishing (or non-negligible) spherical harmonic component.

For a sufficiently dense pixel resolution and sufficiently large FOV, the multipole method will inevitably be more efficient than the direct method. Anticipating that one approach or the other might be more computationally efficient for a given FOV, we implemented both a “direct” algorithm and a “multipole” algorithm, described below.

### A.1. “Direct” algorithm

For the “direct” algorithm, it is important that the sky is divided into patches, or pixels, of equal area. The pixelization scheme defines a surjective map from points on the unit sphere to integer pixel indices,  $f: S^2 \mapsto \mathbb{Z}_{n_{\text{pix}}}$ . It is desirable to be able to compute the discretized version of the rotated FOV,  $b_i(\gamma_i^{-1}\boldsymbol{\omega})$ , by transforming the pixel coordinates themselves. If some pixels are much larger or smaller than others, then some rotations will require a great deal of sub-pixel interpolation. An example of a particularly troublesome pixelization scheme would be a division into an equiangular lattice in right ascension and declination, because pixels near the poles would subtend much smaller solid angles than pixels near the equator. It is much better if the pixels are of equal or nearly equal size, for then a “nearest neighbor” interpolation scheme consists of nothing more than a permutation of the pixel indices. Such an interpolation scheme ought to converge to first order in the pixel area, while being exceptionally fast.

For this reason, we chose to build an implementation of the “direct” algorithm in HEALPix (Górski et al. 2005), a special pixelization of the sphere and an accompanying software package that has become one of the workhorses of all-sky astronomical data analysis. HEALPix stands for Hierarchical Equal Area isoLatitude Pixelization. As the name says, it has three chief virtues. The first is that successive levels of resolution are derived by subdividing pixels. This makes it straightforward to downsample (decreasing resolution) or upsample (increase resolution) a map. Also, for search operations it provides a log-linear speedup that is analogous to binary searches trees or octrees. The second virtue is that all HEALPix pixels at a given resolution have the same area, satisfying the requirement for our “direct” convolution algorithm. Finally, the isolatitude property, that pixels are arranged on rings of constant latitude, facilitates computationally efficient multipole transforms at any resolution, which we exploit in the next section.

---

**Algorithm 1** “Direct” convolution for single telescope case.

---

```

{ $s_i[0], s_i[1], \dots, s_i[n_{\text{pix}} - 1]$ }  $\leftarrow$  pixelized map
{ $v_i[0], v_i[1], \dots, v_i[n_{\text{pix}} - 1]$ }  $\leftarrow$  pixelized mask
{ $b_i[0], b_i[1], \dots, b_i[n_{\text{pix}} - 1]$ }  $\leftarrow$  pixelized FOV
{ $p_i[0], p_i[1], \dots, p_i[n_{\text{pix}} - 1]$ }  $\leftarrow$  {0, 0,  $\dots$ , 0}
for  $j = 0 \rightarrow n_{\text{pix}} - 1$  do
   $\gamma \leftarrow \text{pix2ang}(j)$ 
  for  $k \in \mathbb{Z}: 0 \leq k < n_{\text{pix}}, b_i[k] \neq 0$  do
     $\omega \leftarrow \text{pix2ang}(k)$ 
     $k' \leftarrow \text{ang2pix}(R[\gamma^{-1}]\omega)$ 
     $p_i[j] \leftarrow p_i[j] + b_i[k] \times v_i[k'] \times s_i[k']$ 
  end for
end for

```

---

## A.2. “Multipole” algorithm

The “multipole” algorithm is the analog of FFT convolution in spherical coordinates. Our implementation of fast convolution with HEALPix is based on [Wandelt & Górski \(2001\)](#). A weighted outer product is formed from the coefficients of the spherical harmonic expansions of the sky map and the telescope FOV. The resultant tensor is the 2D discrete Fourier transform (DFT) of the convolution on a equiangular grid in  $(\theta, \phi)$ . Our contribution is to specify a computationally efficient procedure for obtaining the convolution on the HEALPix grid instead of the equiangular grid in spherical coordinates. This is an especially useful feature for our application because output in HEALPix coordinates facilitates further spatial manipulation after the convolution. A detailed explanation of the algorithm is provided below. Source code in C++ is provided in the supplementary material.

Without loss of generality, from here on we drop the environmental factor  $v_i(\omega)$  and the telescope index  $i$ , leaving just

$$p(\text{EM}|\gamma, \text{GW}) = \int s(\omega) b(R^{-1}[\gamma]\omega) d\Omega.$$

Though the integral involves a *passive* rotation of the beam by  $R^{-1}[\gamma]$ , it is conceptually simpler to think of this is as an *active* rotation by  $R[\gamma]$ . In the language of spherical harmonics, this integral can be evaluated at any  $\gamma = (\theta, \phi)$  through the sum

$$p(\text{EM}|\theta, \phi, \text{GW}) = \sum_{l=0}^L \sum_{m=-l}^l \sum_{n=-l}^l s_{lm}^* D_{mn}^l(\phi, \theta, 0) b_{ln} \quad (\text{A1})$$

where

$$D_{mn}^l(\alpha, \beta, \gamma) = e^{-im\alpha} d_{mn}^l(\beta) e^{-in\gamma} \quad (\text{A2})$$

is the Wigner matrix element denoting the active rotation first by  $\gamma$  about the  $z$ -axis, then by  $\beta$  about the new  $y$ -axis, and finally about  $\alpha$  about the final  $z$ -axis. ([Sakurai & Tuan \(1994\)](#) use the active convention; [Edmonds \(1957\)](#) and *Mathematica* employ the passive convention.)  $s_{lm}$  and  $b_{ln}$  are the spherical harmonic coefficients of  $s(\omega)$  and  $b(\omega)$ , respectively, assumed to be zero or at least negligibly small for  $l > L$ .  $d_{mn}^l(\beta)$  is the reduced Wigner matrix element representing an active rotation about the  $y$ -axis by  $\beta$ . The basic idea is to employ a Fourier expansion of the reduced Wigner coefficients ([Edmonds 1957](#); [Risbo 1996](#); [Wandelt & Górski 2001](#)),

$$d_{mn}^l(\beta) = i^{m-n} \sum_{m'=-l}^l \Delta_{m'm}^l \Delta_{m'n}^l e^{-im\beta}, \quad (\text{A3})$$

where  $\Delta_{mn}^l = d_{mn}^l(\pi/2)$ , to convert the summation in Equation (A1) to a two-dimensional, discrete Fourier transform,

$$\begin{aligned} p(\text{EM}|\theta, \phi, \text{GW}) &= \sum_{l=0}^L \sum_{m=-l}^l \sum_{m'=-l}^l \sum_{n=-l}^l i^{m-n} s_{lm}^* \Delta_{m'm}^l \Delta_{m'n}^l e^{-im'\theta} e^{-im\phi} b_{ln} \\ &\equiv \sum_{m=-L}^L \sum_{m'=-L}^L T_{mm'} e^{-im'\theta} e^{-im\phi} \end{aligned} \quad (\text{A4})$$

First, compute both  $s_{lm}$  and  $b_{lm}$  using the HEALPix routine `map2alm` for  $m \geq 0$ . Since  $s(\omega)$  and  $b(\omega)$  are purely real, the remaining coefficients  $m < 0$  are determined by conjugate symmetry,  $s_{l,-m} = (-1)^m s_{lm}^*$ . Also, all of the  $m = 0$  coefficients are purely real.

Next, compute  $\bar{s}_{lm} = (-i)^m s_{lm}$  and  $\bar{b}_{lm} = (-i)^m b_{lm}$  for  $m \geq 0$  to transfer the phase factor  $i^{m-n}$  to the beam and sky map.

Next, evaluate the sum over  $n$  in  $\mathcal{O}(L^3)$  time, yielding  $X_{lm'}$ :

$$X_{lm'} \equiv \sum_{n=-l}^l \Delta_{m'n}^l \bar{b}_{ln} = \Delta_{m'0}^l b_{l0} + 2 \sum_{m=1}^l \Delta_{m'n}^l \begin{cases} \text{Re}[\bar{b}_{ln}] & \text{if } (l-m-n) \text{ is even} \\ \text{Im}[\bar{b}_{ln}] & \text{if } (l-m-n) \text{ is odd} \end{cases} \quad (\text{A5})$$

The expression after the equality sign makes use of  $\bar{b}_{ln} = \bar{b}_{l,-n}^*$  and the identity (McEwen & Wiaux 2011)  $d_{mn}^l(\pi - \beta) = (-1)^{l-m} d_{m,-n}^l(\beta) \implies \Delta_{mn}^l = (-1)^{l-m} \Delta_{m,-n}^l$ .

Next, for  $0 \leq m \leq L$  and  $-L \leq m' \leq L$ , the sum over  $l$  is computed, also in  $\mathcal{O}(L^3)$  time:

$$T_{mm'} \equiv \sum_{l=\max(|m|, |m'|)}^L \bar{s}_{lm}^* \Delta_{m'm}^l X_{lm'} = \sum_{l=\max(m, |m'|)}^L \bar{s}_{lm}^* \Delta_{|m'|m}^l \begin{cases} X_{lm'} & \text{if } m' \geq 0 \\ X_{l|m'|}^* (-1)^{l-m-m'} & \text{if } m' < 0 \end{cases}. \quad (\text{A6})$$

At this point, it is possible to recover the convolution on a equiangular grid on right ascension and declination by performing a 2D forward Fourier transform of  $T_{mm'}$ . However, this grid will oversample the poles; it is preferable to obtain the convolution at the HEALPix pixel sites. The HEALPix grid is arranged such that the pixel centers  $(\theta_j, \phi_{jk})$  lie on a set of  $\mathcal{O}(L)$  iso(co)latitude rings at colatitudes  $\theta_j$ . For the colatitude  $\theta_j$  particular to each ring, the sum over  $m'$  is performed in a *total* of  $\mathcal{O}(L^3)$  operations, yielding

$$T_m(\theta_j) \equiv \sum_{m'=-L}^L T_{mm'} e^{-im'\theta_j}. \quad (\text{A7})$$

This 1D Fourier series,  $T_m(\theta_j)$ , is now shifted in phase so that  $T_0(\theta_j)$  is the value at the first pixel in the ring centered at  $(\phi_{j0}, \theta_j)$ , resulting in

$$T'_m(\theta_j) \equiv T_m(\theta_j) e^{-im\phi_{j0}}. \quad (\text{A8})$$

Since the convolution must be everywhere real, we know that  $T'_{-m}(\theta_j) = [T'_m(\theta_j)]^*$ .

Now, in rings close to the poles, the number of pixels per ring  $n_{\text{ring},j}$  is small; we assume that in *all* rings  $n_{\text{ring},j} < 2L + 1$ . For more efficient evaluation, the Fourier series can be folded, or aliased, to  $T''_m(\theta_j)$ , for  $0 \leq m < n_{\text{ring},j}/2$ , such that

$$p(\text{EM}|\theta_j, \phi_{jk}, \text{GW}) = \sum_{m=-L}^L T'_m(\theta_j) \exp\left[-\frac{2\pi i k m}{n_{\text{ring},j}}\right] = \sum_{m=-(n_{\text{ring},j}/2-1)}^{n_{\text{ring},j}/2-1} T''_m(\theta_j) \exp\left[-\frac{2\pi i k m}{n_{\text{ring},j}}\right] \quad (\text{A9})$$



where

$$T_m''(\theta_j) = \sum_i \left( T_{m+n_{\text{ring},j}i}' + T_{m-n_{\text{ring},j}i}' \right) \text{ for } 0 \leq m < n_{\text{ring},j}/2. \quad (\text{A10})$$

The Fourier series can be summed directly in a total of  $\mathcal{O}(L^3)$  operations for all rings. Although not necessary to preserve the overall cubic scaling with  $L$ , this final step can be sped up by employing a FFT instead of a direct sum. FFT packages such as `fftw`<sup>7</sup> conventionally provide a *real-to-complex* forward FFT for real-valued input data, producing just the nonnegative frequency components, and a corresponding inverse *complex-to-real* FFT converting just nonnegative frequency components back to real data. We would like to use a *complex-to-real* transform to fill in the ring pixels from the Fourier series, but the negative sign in the exponent of Equation (A9) makes it a forward, rather than an inverse, transform. Application of a *complex-to-real* FFT would result in the ring pixels coming out in reversed order. For this reason, we compute

$$T_m'''(\theta_j) = [T_m''(\theta_j)]^* \quad (\text{A11})$$

and then perform the *complex-to-real* FFT on *this* to fill in the pixels in each ring.

### A.3. Comparison

We implemented both the “direct” and “multipole” convolution algorithms in C/C++. The outermost loops are accelerated using OpenMP<sup>8</sup>, a widely supported compiler extension for generating parallel code for multiprocessor machines. Much of the HEALPix library itself, including the spherical harmonic transform routine `map2alm`, is also accelerated with OpenMP. For the multipole algorithm, the final DFTs are performed using `fftw`.

Operating with one thread on a quad core Sun Fire X2200 M2 and a pixel resolution of  $\approx 0.05 \text{ deg}^2$ , the direct algorithm ran in 8.9 s for the ROTSE III FOV, 38.4 s for QUEST, and 931.2 s for Pi of the Sky South. The multipole algorithm, whose run time does not depend on the FOV, took 23.4 s. Using 8 threads, ROTSE-III took 1.3 s, QUEST 5.4 s, Pi of the Sky South 128.9 s, and the multipole algorithm took 6.6 s.

---

<sup>7</sup><http://www.fftw.org/>

<sup>8</sup><http://openmp.org/>

## REFERENCES

- Aasi, J., Abadie, J., Abadie, B. P., & et al. 2012a, in preparation
- . 2012b, LIGO-P1100038-v4, in preparation
- Abadie, J., et al. 2010, *Class. Quantum Grav.*, 27, 173001
- . 2011, *ApJ*, 734, L35
- . 2012a, arXiv:1112.6005v4
- . 2012b, *A&A*, 540, A124
- Abbott, B., et al. 2008, *Phys. Rev. Lett.*, 101, 211102
- Abbott, B. P., et al. 2009, *Rep. Prog. Phys.*, 72, 076901
- Acernese, F., et al. 2008, *Class. Quantum Grav.*, 25, 184001
- Akerlof, C. W., et al. 2003, *PASP*, 115, 132
- Babul, A., Paczynski, B., & Spergel, D. 1987, *ApJ*, 316, L49
- Baltay, C., et al. 2007, *PASP*, 119, 1278
- Barron, E. G., Kaplan, G. H., Bangert, J., Bartlett, J. L., Puatua, W., Harris, W., & Barrett, P. 2011, in *BAAS*, Vol. 43, 217th Meeting of the American Astronomical Society, #344.14
- Belczynski, K., Perna, R., Bulik, T., Kalogera, V., Ivanova, N., & Lamb, D. Q. 2006, *ApJ*, 648, 1110
- Bennett, M. F., van Eysden, C. A., & Melatos, A. 2010, *MNRAS*, 409, 1705
- Berezinsky, V., Hnatyk, B., & Vilenkin, A. 2001, *Phys. Rev. D*, 64, 043004
- Boër, M., et al. 2003, *ESO Messenger*, 113, 45
- Cannon, K., et al. 2012, *ApJ*, 748, 136
- Christensen, N., & Meyer, R. 1998, *Phys. Rev. D*, 58, 082001
- Corsi, A., & Owen, B. J. 2011, *Phys. Rev. D*, 83, 104014
- Coward, D. M., et al. 2010, *PASA*, 27
- Damour, T., & Vilenkin, A. 2000, *Phys. Rev. Lett.*, 85, 3761
- Edmonds, A. 1957, *Angular Momentum in Quantum Mechanics* (Princeton University Press)
- Fairhurst, S. 2009, *New J. Phys.*, 11, 123006
- . 2011, *Class. Quantum Grav.*, 28, 105021
- Gomboc, A., et al. 2004, *Nucl. Phys. B Proc. Supp.*, 132, 312
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759
- Guidorzi, C., et al. 2006, *PASP*, 118, 288
- Harry, G. M., & the LIGO Scientific Collaboration. 2010, *Class. Quantum Grav.*, 27, 084006
- Ingber, L. 1989, *Mathematical and Computer Modelling*, 12, 967
- Ioka, K. 2001, *MNRAS*, 327, 639
- Kanner, J., Huard, T. L., Márka, S., Murphy, D. C., Piscionere, J., Reed, M., & Shawhan, P. 2008, *Class. Quantum Grav.*, 25, 184034
- Keller, S. C., et al. 2007, *PASA*, 24, 1
- Klimenko, S., et al. 2011, *Phys. Rev. D*, 83, 102001
- Klotz, A., et al. 2008, *PASP*, 120, 1298
- Kocsis, B., Haiman, Z., & Menou, K. 2008, *ApJ*, 684, 870
- Kuroda, K. 2010, *Class. Quantum Grav.*, 27, 084004
- Law, N. M., et al. 2009, *PASP*, 121, 1395
- Lee, W. H., & Ramirez-Ruiz, E. 2007a, *New J. Phys.*, 9, 17
- . 2007b, *New J. Phys.*, 7
- Majcher, A., et al. 2011, *Proc. SPIE*, 8008
- Malek, K., et al. 2010, *Adv. Astron.*, 2010, 9
- Manzotti, A., & Dietz, A. 2012, arXiv:1202.4031
- McEwen, J., & Wiaux, Y. 2011, *IEEE Trans. Signal Process.*, 59, 5876

- Mészáros, P. 2006, Rep. Prog. Phys., 69, 2259
- Metzger, B. D., & Berger, E. 2012, ApJ, 746, 48
- Nakar, E. 2007, Phys. Rep., 442, 166
- Nissanke, S., Holz, D. E., Hughes, S. A., Dalal, N., & Sievers, J. L. 2010, ApJ, 725, 496
- Nissanke, S., Sievers, J., Dalal, N., & Holz, D. 2011, ApJ, 739, 99
- Ott, C. D. 2009, Class. Quantum Grav., 26, 063001
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, Numerical Recipes, 3rd edn. (Cambridge University Press), 641–647
- Raymond, V., van der Sluys, M. V., Mandel, I., Kalogera, V., Rver, C., & Christensen, N. 2009, Class. Quantum Grav., 26, 114007
- Risbo, T. 1996, J. Geodesy, 70, 383, 10.1007/BF01090814
- Sakurai, J., & Tuan, S. 1994, Modern quantum mechanics (Addison-Wesley Pub. Co.)
- Sathyaprakash, B., & Schutz, B. F. 2009, Living Rev. Relativ., 12
- Schutz, B. F. 2001, in Proceedings of the MPA/ESO/MPE/USM Joint Astronomy Conference, ed. M. Gilfanov, R. Sunyaev, & E. Churazov, ESO Astrophysics Symposia, Garching, Germany
- Schutz, B. F. 2011, Class. Quantum Grav., 28, 125023
- Searle, A. C., Sutton, P. J., & Tinto, M. 2009, Class. Quantum Grav., 26, 155017
- Searle, A. C., Sutton, P. J., Tinto, M., & Woan, G. 2008, Class. Quantum Grav., 25, 114038
- Siemens, X., Creighton, J., Maor, I., Majumder, S. R., Cannon, K., & Read, J. 2006, Phys. Rev. D, 73, 105001
- Sylvestre, J. 2003, ApJ, 591, 1152
- Troja, E., King, A. R., O’Brien, P. T., Lyons, N., & Cusumano, G. 2008, MNRAS, 385, L10
- U.S. Nautical Almanac Office. 2010, Astronomical Almanac: 2011, ASTRONOMICAL ALMANAC FOR THE YEAR (Bernan Assoc)
- Vachaspati, T. 2008, Phys. Rev. Lett., 101, 141301
- van der Sluys, M., Raymond, V., Mandel, I., Rver, C., Christensen, N., Kalogera, V., Meyer, R., & Vecchio, A. 2008, Class. Quantum Grav., 25, 184011
- Veitch, J., & Vecchio, A. 2010, Phys. Rev. D, 81, 062003
- Virgo Collaboration. 2009, VIR-027A-09, <https://tds.ego-gw.it/ql/?c=6589>
- Wandelt, B. D., & Górski, K. M. 2001, Phys. Rev. D, 63, 123002
- Wang, Y., Fan, Y.-Z., & Wei, D.-M. 2011, Phys. Rev. Lett., 106, 259001
- Wen, L., & Chen, Y. 2010, Phys. Rev. D, 81, 082001
- White, D. 2012, personal communication